

Claim Listing

1. (Previously presented) A Programmable Streaming Data Processor (PSDP), arranged to perform primitive initial processing functions directly on a set of data comprising:
 - a streaming data interface, arranged to receive non-field delineated data from a streaming data source;
 - a streaming interface First In First Out (FIFO), arranged to temporarily store the streaming non-field delineated data from the streaming data interface;
 - a data engine, arranged to receive the non-field delineated output data from the streaming interface FIFO, recognize the record and field structure of the non-field delineated data, determine field boundaries in the non-field delineated data, and process fields to select one or more fields to be assembled into output tuples, the data engine also containing logic arranged to determine whether an output tuple is to be selected for further processing by additional Job Processing Units (JPUs) and to assert a use or lose decision value according to that determination;
 - a tuple generator, arranged to assemble fields into the output tuple and, if the use or lose decision value indicates that such output tuple is to be discarded, to prevent the output tuple from being transferred for further processing by the JPU; and
 - an output FIFO device, arranged to temporarily store tuples prior to conditionally forwarding them to the JPU.
2. (Previously presented) An apparatus as in claim 1 wherein the use or lose decision value indicates a result from logic processing of fields read from the streaming data interface.
3. (Previously presented) An apparatus as in claim 1 wherein the use or lose decision value indicates a result from Transaction Identifier (TID) processing.
4. (Original) An apparatus as in claim 3 wherein the TID processing and data engine logic execute in parallel.

5. (Previously presented) An apparatus as in claim 1 wherein the output tuple is greater in length than an expected predetermined size, and the use or lose decision value is then used to set an overflow field in the output tuple.
6. (Previously presented) An apparatus as in claim 5 wherein the use or lose decision value is not asserted when a buffer local to the programmable data streaming processor is full; and
means for appending an overflow filter bit to a tuple that indicates a transfer of a tuple that should be ignored.
7. (Previously presented) An apparatus as in claim 1 additionally comprising:
a Direct Memory Access (DMA) interface, coupled to the output FIFO, to provide direct access to a memory in the JPU.
8. (Previously presented) An apparatus as in claim 1 wherein the use or lose decision value is used to reset an output FIFO write pointer so any prior fields in the output tuple are discarded.
9. (Previously presented) An apparatus as in claim 1 wherein an overflow filter bit is inserted in a length field appended to the output tuple.
10. (Previously presented) An apparatus as in claim 1 wherein an invalid field is appended to a tuple to indicate the results of TID processing.
11. (Previously presented) An apparatus as in claim 10 wherein the results of TID processing indicate that a tuple is to be returned.
12. (Original) An apparatus as in claim 10 wherein the invalid field indicates that the tuple should not have been returned but the output FIFO overflowed.

13. (Previously presented) An apparatus as in claim 1 further comprising:
a register reflecting the final PSDP status which is read by a Central Processing Unit (CPU) to identify whether any overflow or TID status bits are set in any of the tuples.
14. (Previously presented) An apparatus as in claim 1 wherein the use or lose decision value represents DeMorgan's Law reduction of multiple instructions.